

Remarks

Status of application

Claims 1-25 are pending in the current application. The claims stand rejected in view of cited prior art. Claim 7 has been amended to correct a minor typographical error. In view of clarifying remarks made below, re-examination and reconsideration are respectfully requested.

The invention

An e-mail system constructed in accordance with the present invention includes a composer module ("Composer"), a message transport agent (MTA), and a mass-mail accelerator (MMA). The Composer (which itself is a conventional component) is a program that operates against a very large database of users to provide large-scale customized e-mail messages by combining different pieces of a message together on a per-user basis. Ordinarily, the Composer passes a given message on to an MTA that, in turn, transmits the message to the intended recipient. However, in accordance with the present invention, the basic operation is modified so that the Composer passes a given message on to the MMA, which serves to carry out e-mail delivery/routing for the messages that have been passed on to it. More particularly, the degree of parallelism on the MTA side of message delivery has been greatly increased.

In operation, the MMA receives input that, in turn, is fed into one or more queues. The input that is received, via SMTP, comprises outgoing messages from one or more Composers. A receiving (or "client") thread initially handles this input. In the instance that multiple Composers are connected to the MMA (i.e., multiple concurrent connections), one client thread is assigned to each incoming connection. Two types of threads are actually employed here: a "listener" thread waits for a new connection, creates a client thread, and assigns the new connection to it (and thereafter repeats), and a "client" thread is what actually interacts with the Composer beyond the initial TCP/IP handshake. The respective client thread receives the incoming e-mail message (or simply, "message") and, in turn, decides which queue from the set of queues within the MMA is appropriate to receive and process the message. Any number of queues may be supported, as desired

(and as indicated by the ellipsis). The client thread that receives the message examines the configuration and state of the available queues to see which one is appropriate to receive the incoming message.

Each queue itself owns a thread that manages a list of messages. As a particular advantage, the queues themselves are configurable to either be general (generic) or be specific to a particular mail (destination) domain. For instance, a queue may be configured to handle only mail destined for the Hotmail.com domain, or configured to handle only mail destined for the AOL.com domain. A queue that is specifically configured will only handle e-mail for its specific domain and will not handle any other e-mail. In contrast, a queue may be configured to be generic or general, in which case it will handle e-mail destined for any domain which has no specific queue assigned to it. E-mail posted to a specific queue will not require a Domain Name Services (DNS) look-up, as the MMA already knows (i.e., has cached) the DNS information for the corresponding target e-mail domain. Thus, for example, e-mail destined for the AOL.com domain is posted to the AOL queue. The MMA need not look up the DNS information for the AOL.com domain as this information has already been cached as part of the setup for the AOL queue. Which queues are created is entirely dependent on the configuration which gives the customer-user (e.g., system administrator) the ability to tailor or tune for a given situation. If, for example, the system administrator knows that about 60% of outgoing e-mail for his or her company is going to AOL, then the system administrator would set up an AOL-specific queue, with corresponding resources.

Each queue manages a pool of MTA threads. During configuration of the queues, the customer's system administrator may specify the allocation of MTA threads to a given queue. For instance, a system administrator may specify a maximum and/or minimum number of MTA threads that are available to a given queue. When a given MTA thread is started, it establishes a connection out to a real MTA (e.g., remote MTA residing at a particular destination on the Internet). This connection is established using SMTP over a TCP (Transmission Control Protocol) connection. Via this connection, a given MTA thread may talk SMTP to an actual MTA out in the real world someplace (e.g., an AOL MTA).

During MMA operation, once a message has been passed to a queue, that queue examines its MTA threads to see if one is ready to accept the message. If an MTA thread is ready, the queue will assign the message to that MTA. Once a message is assigned to an MTA thread, that thread is no longer available and, thus, it marks itself as "busy" (or otherwise removes itself from a "ready" list). The MTA thread proceeds to handle the work of the SMTP exchange between the MMA and the target real-world MTA (e.g., AOL MTA). While a given MTA thread is waiting for a reply from the destination MTA (e.g., AOL MTA), the MMA can proceed to do other work. Thus, for instance, while a given message is being handled by a particular MTA thread, other incoming messages can be injected, queued, requeued, moved around, or the like, within the system. In this manner, the bottleneck usually encountered with processing mass e-mailings is removed.

Prior art rejections

A. Section 103: Sriram (US Patent No. 5,463,620)

Claims 1-7 and 13-24 stand rejected under 35 U.S.C. 103(a) as being unpatentable over consideration of US Patent 5,463,620 to Sriram. Regarding independent claims 1, 13, and 21, for example, the Examiner states the following basis for rejection:

establishing a plurality of queues in the system, zero or more of these being specific queues for handling mail to a specific set of domains, and one being a general queue for transferring e-mail to domains not handled by specific queues, (Cols. 11-16; note Col. 12, lines 4-53);

receiving at the system a request to process for transfer a plurality of outbound e-mail message (threads), each e-mail message specifying delivery to at least one recipient at a particular domain, (Figs. 1 & 5; Col. 4, lines 43-67; Col. 4; and Col. 5, lines 1-34); and

for each given e-mail message, processing the given e-mail message by:

determining what domain the given e-mail message is destined for, if the determined domain for the given e-mail message is a specific domain handled by a corresponding specific queue, assigning the given e-mail message to the corresponding specific queue for transferring the given e-mail to said specific domain, otherwise assigning the given e-mail message to said general queue, and without waiting for confirmation that the given e-mail message has been

successfully processed for transfer to another system, proceeding to process the next one of the e-mail messages, (Col. 11, lines 2-16).

The Examiner goes on to say that "within a queuing system, like that of Sriram, confirmation would not be necessary, as it would defeat the purpose of efficiently integrating diverse traffic types on a single network." A detailed review of the Sriram patent reveals that Applicant's claimed invention may be distinguished on a variety of grounds. And, in fact, Sriram has little relevancy to Applicant's invention.

Sriram describes a queuing system for improving transmission and bandwidth allocation in broadband Asynchronous Transfer Mode (ATM) networks. This pertains to a low-level transport protocol for routing packets over a specific medium. In terms of the OSI (Open System Interconnection) networking model, Sriram's invention applies at the lowest layers such as the Physical, Data-Link, or Network Layers. Software operating at those layers is completely unaware (and rightly so) of the nature of the data it is routing. Accordingly, such software is hardly in the position to intelligently act upon data which the software itself (at this low-level layer in the OSI model) has no facility to interpret or act upon. (For the Examiner's convenience, a brief summary of the OSI model is included as Attachment #1.) In fact, the only reference to "e-mail" anywhere in Sriram is a vague passing reference that the many types of traffic that could be improved by Sriram's improved routing approach could also include (of course) e-mail traffic. Processes operating at low-level layers in the OSI model cannot intelligently act on the data payload embedded in the transmission packets without violating the requirements of the OSI model.

In contrast to the low-level message processing in Sriram, all e-mail software making connection decisions operates at the Application, Presentation, and Session Layers of the OSI model. TCP/IP operates at the Transport and Network layers. Applicant's MMA invention is an application and as such operates at only the highest three layers. It makes calls to Berkeley UNIX TCP functions such as socket(), listen(), bind(), accept() and connect(), which are its interface to the Transport and Network layers. It has no direct interaction with the lower layers of the protocol stack. Sriram's invention operates inside routers or switches, while Applicant's invention runs in

computer servers and is in fact unaware of the software running in the switches it uses to connect to the Internet.

The configuration of Applicant's MMA invention allows an administrator to note that, for example, most of his or her traffic goes to a specific few domains, and as a result dedicate certain memory and socket-level resources (all still at the higher layers) to prepare for traffic going to those destinations. Here, a "queue" is a set of resources used to send bulk mail to a specific domain or set of servers. In contrast, Sriram's system allocates into distinct queues various types of ATM cells (similar to IP packets) rather than those going to common destinations, and Sriram describes algorithms for fairly servicing those queues. In Sriram, a "queue" is merely a list of packets of common priority waiting to get serviced (i.e., written out to the network medium for transport). The two are fairly unrelated concepts other than common use of the word "queue".

Given the above noted differences, it is no surprise to find that the Sriram patent is entirely silent regarding the dedication of resources to routing traffic to specific destinations known to be "busy", either in general or with respect to a specific application or origin. Quite simply, Sriram's approach does not focus on the destination of any of the traffic it seeks to route.

Turning now to Applicant's claims, one finds specific limitations that are left unaddressed by Sriram. Applicant does not claim to have invented some general notion about how message traffic may be queued or routed. Instead, Applicant has invented a specific solution for processing mass mailing (i.e., bulk) e-mail messages that are being sent to recipients at various destination domains. These distinctions are set forth in particular detail in Applicant's claims. For example, claim 1 includes the following claim limitations (among others):

for each given e-mail message, processing the given e-mail message by:

determining what domain the given e-mail message is destined for,

if the determined domain for the given e-mail message is a specific domain handled by a corresponding specific queue, assigning the given e-mail message to the corresponding specific queue for transferring the given e-mail to said specific domain, otherwise assigning the given e-mail message to said

general queue, and

without waiting for confirmation that the given e-mail message has been successfully processed for transfer to another system, proceeding to process the next one of the e-mail messages.

As shown, the claim requires determining what domain a given e-mail message is destined for. Sriram has no facility to perform such a determination (and making a determination at such a low level would be in violation of the OSI model). For example, Sriram states that his invention entails "classifying each call in accordance with certain signal characteristics, such as required bandwidth and sensitivity to delay. Each call class is directed to a separate queuing circuit." (See, e.g., Sriram Abstract.) Note in particular that Sriram's approach is not queuing traffic based on the determined destination of that traffic, as required by Applicant's claim limitations. To contend that Sriram's approach teaches Applicant's claimed approach of queuing e-mail messages based on destination domains ignores the most basic teaching of Sriram.

Applicant's other rejected independent claims 13 and 21 include similar limitations. For example, claim 13 requires:

a processing thread for receiving incoming e-mail messages that are to be transferred to another system, and **assigning each incoming e-mail message to a particular queue based on what domain the incoming e-mail message is destined for;**

(emphasis added)

Similarly, claim 21 recites:

dividing incoming e-mail messages that are to be processed for transfer into different groups, **based on what domain each e-mail message is destined for;**

(emphasis added)

As shown, each of Applicant's independent claims includes claim limitations requiring that the assignment of e-mail messages into different queues be based on a determination

of the particular **destination domain** that each e-mail message is intended for. Sriram does not teach or suggest such a feature. Instead, if one were to apply the teachings of Sriram to an e-mail system, one would have an e-mail system in which the message traffic is prioritized at the level of the router to better address required bandwidth and sensitivity to delay. One would not have, however, a solution to the bottleneck that accompanies processing mass e-mailings, and therefore the opportunity for significant improvements in scalability and throughput would not be realized.

The rejected dependent claims are likewise believed to be allowable for the reasons stated above. Further, the dependent claims include additional limitations that further distinguish Applicant's invention. For example, dependent claims 2-4 set forth the claim limitation of a general queue and one or more optional specific queues. The Examiner states that this is taught at Col. 10, lines 58-67. However, that section merely states:

Summary

This example of the invention provides a way to conveniently handle and integrate a wide variety of communications traffic handled by a broadband network. Bandwidth is effectively and fairly allocated and congestion is avoided. The dynamic time slice server guarantees a desired bandwidth to calls requiring a fixed wide bandwidth for the duration of the call which facilitates setting up circuit-like constant bit rate connections in an asynchronous transfer mode network. Statistical multiplexing of calls is done when applicable. Call admittance is based upon capacity versus [...]

A review of the above passage reveals no discussion of general versus specific queues, or any feature of Sriram that would be analogous. The Examiner also points generally to Cols. 11-16. However, the text of Sriram at Cols. 11-16 mainly consists of Sriram's claims. There, Sriram describes various special-purpose queues, such as one for low bandwidth, or for high bandwidth delay sensitive calls, or for high bandwidth delay insensitive calls, etc. There is no discussion of a general or "catch all" queue that would operate in the manner of Applicant's general queue (i.e., for default queuing of traffic that is not destined for one of the specific queues).

In summary, Sriram is an ATM traffic sorting and prioritizing algorithm, which

seeks to optimize use of an ATM network, with particular focus on the switches/routers in which it operates. To the point, Sriram's system does not take into account any previously-known information about the data (i.e., data payload) it seeks to transport. Applicant's MMA invention is also a traffic sorting and prioritizing algorithm, but does have prior knowledge about its data set, including scope of operation outside of the computer on which it is operating; it has no overlap in terms of the protocol layers in use (in contrast to Sriram). Given these shortcomings of Sriram, it is respectfully submitted that Applicant's claimed invention distinguishes over Sriram and is patentable under Section 103.

B. Section 103: Sriram in view of Vaid

Claims 8-12 and 25 stand rejected under 35 U.S.C. 103(a) as being unpatentable over US Patent 5,463,620 to Sriram in view of US Patent US 6,502,131 B1 to Vaid. Here, the Examiner relies on Sriram for the base rejection, but adds Vaid for the "composer" limitation of the claims. The claims (which depend from independent claims 1, 13, and 21) are believed to be allowable for at least the reasons stated above pertaining to Sriram. The claims are also believed to be allowable for the following additional reasons.

The Examiner generally cites Vaid's Cols. 14-17, and specifically cites the following section (Col. 7, lines 51-66):

As shown above, there exist a large number of diverse applications and protocols that are widely used and have their own performance requirements. For example, applications such as mail (e.g., SMTP) and news (e.g., NNTP) are not interactive and are therefore not sensitive to delay. On the other hand, applications such as real-time conferencing are extremely sensitive to delay but not to packet loss. Applications such as TELNET and DNS do not utilize significant bandwidth, but are sensitive to delay and loss. Conversely, applications such as FTP consume a great deal of bandwidth but are not that sensitive to delay. Generally, network applications can be categorized as:

1. Interactive (e.g., delay sensitive) versus non-interactive (e.g., delay tolerant);

2. Bandwidth intensive (bulk data) versus non-bandwidth intensive; and

3. Bursty versus non-bursty.

Here, the Examiner contends that, "Vaid specifically enumerates the use of SMTP (email) and bulk data, (Col. 7, lines 51-66 & Cols. 14-17), which email programs would obviously include a composition program and which composition program would obviously include a database of information for the sending/receiving/routing of email."

An e-mail "composer" -- by itself -- is a known prior art concept and appears as Applicant's admitted prior art in Applicant's Background section. For example, Applicant's Background section describes in detail the Composer program used by Doubleclick, at page 7 of the Background. Vaid, which is directed to a directory enabled policy management tool for intelligent traffic management, is itself entirely silent on this point (i.e., no discussion of mass e-mailings, no discussion of the use of a database for dynamically customizing e-mails, etc.). Nevertheless, for the reasons stated above a composer program by itself is considered Applicant's admitted prior art. However, Applicant certainly does not contend to have invented the notion of a composer program, and the subject matter of the rejected claims is not drawn to just a composer program itself. Instead, the subject matter of these claims recite Applicant's base method (e.g., claim 1) or system (e.g., claim 21) operating on the automated, high volume output of a composer program.

Vaid itself does not include any teaching that remedies the deficiencies noted above pertaining to Sriram, in regards to Applicant's mass-mail accelerator features. Importantly, Vaid cannot possibly teach or suggest the combination of Applicant's base claimed method (e.g., claim 1) or system (e.g., claim 21) with a composer program, as Vaid itself contains no teaching or suggestion of anything (remotely) related to a composer program that could be combined with Sriram to re-create all of the claim limitations for the rejected claims of this group.

To establish a prima facie case of obviousness under Section 103, the Examiner must establish: (1) that there is some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to

modify the reference or to combine reference teachings, (2) that there is a reasonable expectation of success, and (3) that the prior art reference (or references when combined) must teach or suggest all the claim limitations. (See e.g., MPEP 2142). It is respectfully submitted that the combination of Sriram with Vaid fails to teach or suggest all the claim limitations set forth in the rejected claims of this group, and that the claims are therefore patentable under Section 103.

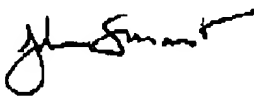
Conclusion

In view of the foregoing remarks and the amendment to the claims, it is believed that all claims are now in condition for allowance. Hence, it is respectfully requested that the application be passed to issue at an early date.

If for any reason the Examiner feels that a telephone conference would in any way expedite prosecution of the subject application, the Examiner is invited to telephone the undersigned at 408 884 1507.

Respectfully submitted,

Date: January 21, 2005

✓  Digitally
signed by
John A. Smart
Date:
2005.01.21
16:13:17
-08'00'

John A. Smart; Reg. No. 34,929
Attorney of Record

408 884 1507
815 572 8299 FAX